

Spur-Reduced Digital Sinusoid Synthesis

Michael J. Flanagan and George A. Zimmerman
Jet Propulsion Laboratory
California Institute of Technology

November 17, 1993

Abstract

This paper presents and analyzes a technique for reducing the spurious signal content in digital sinusoid synthesis. Spur reduction is accomplished through dithering both amplitude and phase values prior to word-length reduction. The analytical approach developed for analog quantization is used to produce new bounds on spur performance in these dithered systems. Amplitude dithering allows output word-length reduction without introducing additional spurs. Effects of periodic dither similar to that produced by a pseudo-noise (PN) generator are analyzed. This phase dithering method provides a spur reduction of $6(M + 1)$ dB per phase bit when the dither consists of M uniform variates. While the spur reduction is at the expense of an increase in system noise, the noise can be made white, making the noise power spectral density small. This technique permits the use of a smaller number of phase bits addressing sinusoid look-up tables, resulting in an exponential decrease in system complexity. Amplitude dithering allows the use of less complicated multipliers and narrower data paths in purely digital applications, as well as the use of coarse-resolution, highly-linear digital-to-analog converters (DACs) to obtain spur performance limited by the DAC linearity rather than its resolution.

1 Introduction

It is well-known that adding a dither signal to a desired signal prior to quantization can render the quantizer error independent of the desired signal [1, 2, 3]. Classic examples of this work deal with the quantization of analog signals. Advances in digital signal processing speed and large scale integration have led to the development of all-digital receiver systems, direct digital frequency synthesizers and direct digital arbitrary waveform synthesizers. In all these applications, since finite word-length effects are a major factor in system complexity, they may ultimately determine whether it is efficient to digitally implement a system with a particular set of specifications. Earlier work [4] has presented a technique for reducing the complexity of digital oscillators through phase dithering with the claim of increased frequency resolution. Recent research [5] has suggested mitigation of finite-word-length effects in the synthesis of oversampled sinusoids through noise shaping. This paper shows how the analysis techniques used for quantization of analog signals can be applied to overcome finite-word-length effects in digital systems. The analysis in this paper shows how appropriate dither signals can be used to reduce word lengths in digital sinusoid synthesis without suffering the normal penalties in spurious signal performance. Furthermore, the dithering technique presented in this paper is not limited to the synthesis of oversampled signals.

Conventional methods of digital sinusoid generation [6], e.g. Fig. 1, result in spurious harmonics (spurs) which are caused by finite word-length representations of both amplitude and phase samples [7]. Because both the phase and amplitude samples are periodic sequences, their finite word-length representations contain periodic error sequences, which cause spurs. The spur signal levels are approximately 6 dB per bit of representation below the desired sinusoidal signal.

The technique presented in this paper reduces the representation word length without increasing spur magnitudes by first adding a low-level random noise, or dither, signal to

the amplitude and/or the phase samples, which are originally expressed in a longer word length. The resulting sum, a dithered phase or amplitude value, is truncated or rounded to the smaller, desired word length. Of course, either the amplitude or the phase or both can be dithered. In phase dithering the spurious response is determined by the type of dithering signal employed. In amplitude dithering the spurious response is determined by the original, longer word length. While the amplitude-related spurious is generally related to the phase-related spurious, we will make the pre-dither amplitude word length long enough to satisfy spur power specifications. Then the exact relationship is unimportant, and since the phase dither signal is independent of the amplitude dither signal, the amplitude and phase dithering processes can be treated independently.

The next section describes the quantizer model. Amplitude and phase quantization effects are reviewed in Sections 3 and 4, and simple new bounds on spurious performance are presented. In contrast to bounds in the existing literature, the new bounds are straightforward and require little information about the signal to be quantized. The derivations of the new bounds provide motivation for new analysis of dithered quantizer performance that occurs later in this paper. An analysis of dithering with a periodic noise source is presented in Section 6. The periodic noise source is considered because of its similarity to implementations involving linear feedback shift registers (LFSRs), or Pseudo-Noise (PN) generators. New analysis of phase dithering effects is presented in Sections 7 and 8, followed by simulation results and a design example.

2 Quantizer Model

When a discrete-time input signal, $x[n]$, is passed through an ideal uniform mid-tread quantizer [8], the output signal, $y[n]$, can always be expressed as $y[n] = x[n] + e[n]$ where $e[n]$ is the quantization error, a deterministic function of $x[n]$. The input to the quantizer is

mapped to one of 2^b levels, where b is the number of bits which digitally represent the input sample. Output levels are separated by one quantizer step size, $\Delta = 2^{-b}$. Throughout this paper Δ_A will be used as the step size for amplitude quantization results, Δ_P will be used for phase quantization results, and Δ will be used if the result applies to both amplitude and phase quantization. Similar subscripting will be used on the quantization error.

The input/output relation of a mid-tread quantizer appears in Fig. 2. If the input does not saturate the quantizer then the quantizer error is [8]:

$$e[n] = \sum_{\substack{k=-\infty \\ k \neq 0}}^{\infty} (-1)^k \frac{\Delta}{j2\pi k} \exp\left(j2\pi k \frac{x[n]}{\Delta}\right). \quad (1)$$

Note that the above equation is not correct when the quantizer input is $\pm \Delta/2 \bmod \Delta$. This is not a problem since future analysis involving this expression will treat the input as having a piecewise continuous probability density function in which case the collection of points in question have probability measure zero. If the input signal is bounded so that $|x[n]| \leq A_Q$ where $A_Q = 1/2 - \Delta$, then the quantizer does not saturate and $|e[n]| \leq \Delta/2$. Throughout this paper, quantizers are always operating in non-saturation mode.

3 Amplitude Quantization Effects

Let a discrete-time sinusoid with amplitude $A \leq A_Q$ and frequency ω_0 be the input to a mid-tread quantizer. If the sinusoid is generated in a synchronous discrete-time system, ω_0 can be expressed as 2π times the ratio of two integers. The input sequence is then periodic. Since the error sequence, $e_A[n]$, is a deterministic function of the input sequence, it is periodic as well. Therefore, the spectrum of the error sequence will consist of discrete frequency components (spurs) which contaminate the spectrum of $x[n]$.

The following argument leads to an upper bound on the size of the largest frequency component in the spectrum of $e_A[n]$. Assuming the quantizer is not saturated by the input

signal $x[n]$, the maximum possible quantization error is $A_A/2$, where A_A is the amplitude quantization step size. The total power in $e_A[n]$ is then bounded by $\Delta_A^2/4$. By Parseval's relation, the sum of the spur powers in the spectrum of $e_A[n]$ equals the power in $e_A[n]$. In order to maximize the power in a given spur, the total number of spurs must be minimized. Since $e_A[n]$ is real, the maximum power in a spur occurs when there are two frequency components at $+\omega_{spur}$ and $-\omega_{spur}$, with equal power¹. With two frequency components the power in a single spur is $\leq \Delta_A^2/8$.

Since $x[n]$ is real, its spectrum consists of a positive and a negative frequency component, each having power $A^2/4$. Using the above bound on spur power, the Spurious-to-Signal Ratio ($SpSR$) is $\leq \Delta_A^2/(2A^2)$. If $A = A_Q \approx 1/2$ provided b is not small, then in decibels with respect to the carrier (dBc), $SpSR \leq 3 - 6b$ dBc, where $A_A = 2^{-b}$, and b is the word length in bits. In summary, this upper bound on power in a spur caused by amplitude quantization exhibits -6 dBc per bit behavior.

4 I²Lbase Quantization Effects

Let a phase waveform, $\phi[n]$, be the input to the fixed-tread quantizer. The phase waveform, $\phi[n] = \langle fn + \Phi/2\pi \rangle$ is a sampled sawtooth with amplitude ranging from 0 to 1. The frequency to be generated, measured in samples/cycle, is f , and the phase, measured in radians, is Φ . The fractional operator, $\langle x \rangle$, is defined so that $\langle x \rangle = x \bmod 1$, e.g., $\langle 1.3 \rangle = 0.3$. Since $\phi[n]$ is generated by a synchronous, finite-word-length, discrete-time system, it has a finite period. The signal output from the quantizer can be expressed as $\phi[n] + e_P[n]$, where $e_P[n]$ represents the error introduced by quantization. Since $\phi[n]$ is periodic, $e_P[n]$ is periodic with a period less than or equal to the period of $\phi[n]$. After multiplication by 2π and passage

¹DC offsets and half sampling rate spurs are excluded because they can be corrected by appropriate calibration and filtering.

through the ideal function generator, the output signal is $y[n] = A \cos(2\pi\phi[n] + 2\pi c_P[n])$. If the quantizer has many levels, i.e., > 16 , $c_P[n] \ll 1$, and the small angle approximation $y[n] \approx A \cos(2\pi\phi[n]) - 2\pi A c_P[n] \sin(2\pi\phi[n])$ may be used.

Since $c_P[n]$ and $\phi[n]$ are periodic, the total error $2\pi A c_P[n] \sin(2\pi\phi[n])$ is periodic. The total error power is bounded by $\pi^2 A^2 \Delta_P^2$, because $c_P[n]$ is bounded by $\Delta_P/2$ and the magnitude of a sinusoid is bounded by unity. Recalling the arguments in the previous section on amplitude quantization effects, the maximum spur power of the real error signal is bounded by placing the total error power into two spectral components. Therefore, the maximum spur power is $\pi^2 A^2 \Delta_P^2/2$, where $\Delta_P = 2^{-b}$ and b bits are used to represent phase samples. By the above approximation for $y[n]$ and the bound on the spur power, the Spurious-to-Signal Ratio bound is $SpSR \leq 2\pi^2 \Delta_P^2 = 13 - 6b$ dBc, independent of the signal amplitude, A . This new, simple derivation demonstrates the underlying -6dBc per phase bit behavior, without the analytical complexity found in other existing bounds. More complicated arguments [7] improve the 1.01111(1) by about 9 (11).

5 Amplitude Dithering

In this section rounding the sum of an already quantized sinusoid and an appropriate dither signal is shown to cause spurious magnitudes which depend on the original (longer) word length, not the output (shorter) word length. This phenomenon occurs at the expense of increased system noise from the addition of the dithering signal. An important finite word-length dithering system is subsequently shown to be equivalent to the continuous-amplitude uniformly-dithered system.

Consider the conceptual block diagram for a waveform generator shown in Fig. 3. The b bit quantizer can be split into two parts as in Fig. 4: a high-resolution B -bit quantizer ($B > b$) followed by truncation or rounding to 1, bits. Thus, the generation process consists

of two separate steps: production of a high-resolution waveform and reduction of the word length. The number of bits used to represent the $(1/g)$ -resolution samples should be sufficient to guarantee the desired spectral purity. Then, the word length should be reduced without creating excess signal-dependent quantization error.

The input in Fig. 5 is a B -bit representation of a sinusoid, $x[n] = A \sin(2\pi\phi[n]) + e_{A0}[n]$, where $e_{A0}[n]$ is the quantization error. The dither signal, $z_u[n]$, is white noise uniformly distributed in $[-\Delta_A/2, \Delta_A/2)$, where $\Delta_A = 2^{-b}$. The sum $z_u[n] + x[n]$ is rounded to retain only the b most significant bits. The rounding can be modeled as a uniform quantizer with step size Δ_A . The amplitude A is chosen to avoid saturating this quantizer when the dither signal is added, i.e., $A + \Delta_A/2 \leq A_Q$.

The output from the quantizer can be expressed as $y[n] = x[n] + z_u[n] + e_A[n]$. The characteristic function of the dither signal, $z_u[n]$, is:

$$P_z(\alpha) = E\{\exp(j\alpha z[n])\} = \frac{2 \sin(\alpha \Delta_A/2)}{\alpha \Delta_A} \text{sinc}\left(\frac{\alpha \Delta_A}{2\pi}\right), \quad (2)$$

which has zeros at non-zero integer multiples of $2\pi/\Delta_A$. Thus, as shown in [1], $e_A[n]$ will be a white, wide-sense stationary process, uniformly distributed over $[-\Delta_A/2, \Delta_A/2)$, and it will not contribute spurious harmonics to the output spectrum of $y[n]$. Any spurious components in $y[n]$ are therefore due to $e_{A0}[n]$, which are present in the B -bit input, $x[n]$.

It remains to comment on the noise power not isolated in discrete spurious frequency components. Gray and Stockham have shown [12] that the power in the signal $y[n] = x[n] + e_A[n] + z_u[n]$ is $\Delta^2/6$. This is approximately twice the error variance of a quantization system with no dithering signal. In summary, $y[n]$, which is quantized to b bits, exhibits spurious performance as if it was quantized to B bits ($B > b$), at the expense of doubling the white noise power.

Because the input $x[n]$ is expressed as a B -bit value, an important system equivalent to continuous-amplitude, uniformly-dithered word-length reduction can be constructed. Re-

place the uniformly distributed dither signal, $z_u[n]$, by a finite word-length representation of it, $z[n]$, which is said to be discretely and evenly distributed over the $(B - b)$ -bit quantized values in the region $[-\Delta_A/2, \Delta_A/2)$. Heuristically, $z[n]$ randomizes the portion of the finite word-length input, $x[n]$, that is about to be thrown away by the rounded truncation. This process is equivalent to continuous uniform dithering, since if $x[n]$ is padded out to an infinite number of bits by placing zeros beyond the least significant bit (LSb), then only the $B - b$ most significant bits of $z_u[n]$ will have an effect on the resulting sum, $x[n] + z_u[n]$. All of the bits below the most significant $B - b$ are added to zero, and cannot beget a carry. The output, $y[n]$, is identical in both systems. Therefore $z_u[n]$, continuously, uniformly distributed over $[-\Delta_A/2, \Delta_A/2)$ can be replaced by the discretely valued $z[n]$, and yield the same spurious response for $y[n]$.

It appears that the finite word-length dither signal, $z[n]$, could be generated by a linear feedback shift register (LFSR), or PN generator. This will be strictly true only if the PN generator has an infinite period, since, at this time, the dither signal is required to be white. However, it is not surprising that ideal behavior is approached as the period of the PN generator gets longer. With a sufficiently long period, the case where spur magnitudes are limited by the original word length can be achieved. The following section gives a simple model for a system implementation using a periodic random sequence which can be approximated by a PN generator.

6 Effect of Periodic Dither

This section analyzes the use of a periodic dither signal with a long period, L , for both amplitude and phase dithering. Since the dither signal is periodic, the discrete frequency components in its spectrum will contaminate the desired signal. It is shown that the period can be chosen to satisfy worst case spurious specifications. In this section, the case where

the dither signal is generated using one uniform variate ($M = 1$) is given. When the dither signal is the sum of M independent uniform variates ($M > 1$), as in Section 8, the analysis is the same because the resulting signal is an i.i.d. sequence of random variables.

Instead of using the white dither process, $z_u[n]$, described in the previous section, consider a substitute, $z_L[n]$, which is periodic with period L . Any two samples, $z_L[n]$ and $z_L[n + m]$, where $m \neq 0 \bmod L$, are independent. Samples of $z_L[n]$ are uniformly distributed between $[-A/2, A/2)$, and the quantization step size is A .

When $z_L[n]$ is used as the dither signal, let the quantizer error be called $e_L[n]$. The autocorrelation of $z_L[n]$ when the lag, m , is all integer multiple of L is equal to $R_{z_L z_L}[0] = A^2/12$. in the PN generator approximation to this noise source, $L = 2^l - 1$ where l is the length of the shift register in bits. At other lag values, the samples of $z_L[n]$ are independent, and since they have zero mean, the autocorrelation is zero. Therefore, expressing the autocorrelation as a discrete-time Fourier series:

$$R_{z_L z_L}[m] = \frac{A^2}{12} \delta[m \bmod L] = \sum_{l=0}^{L-1} \frac{1}{L} \frac{A^2}{2L} \exp(j \frac{2\pi m l}{L})$$

where $\delta[m]$ is the Kronecker delta function ($\delta[0] = 1$, $\delta[m] = 0$, $m \neq 0$), and $z_L[n]$ contains L discrete frequency components, each with power $A^2/(12L)$.

In the autocorrelation expression for $e_L[n]$, the expectation is taken over the random variables $z_L[n]$ and $z_L[n + m]$, using the definition of the autocorrelation and Equation 1:

$$R_{e_L e_L}[n, n + m] = \sum_{\substack{k=-\infty \\ k \neq 0}}^{\infty} \sum_{\substack{l=-\infty \\ l \neq 0}}^{\infty} \alpha_k[n] \alpha_l^*[n + m] E \left\{ \exp(j \frac{2\pi}{\Delta} (k z_L[n] - l z_L[n + m])) \right\} \quad (3)$$

where:

$$\alpha_k[n] = \frac{A(-1)^k}{j2\pi k} \exp(j \frac{2\pi k s[n]}{\Delta})$$

The desired signal to which the dither signal $z_L[n]$ is added is $s[n]$. Using the notation from earlier sections, in phase quantization, $s[n] = \phi[n]$ and in amplitude quantization $s[n] = x[n]$.

When the lag is not a non-zero integer multiple of L ,

$$\begin{aligned} E \left\{ \exp\left(\frac{j2\pi}{\Delta}(kz_L[n] - lz_L[n+m])\right) \right\} &= E \left\{ \exp\left(\frac{j2\pi k z_L[n]}{\Delta}\right) \right\} E \left\{ \exp\left(\frac{-j2\pi l z_L[n+m]}{\Delta}\right) \right\} \\ &= E_z\left(\frac{2\pi k}{\Delta}\right) E_z\left(\frac{-2\pi l}{\Delta}\right) = \delta[k]\delta[l]. \end{aligned}$$

This last fact is true because the characteristic function of $z_L[n]$ has zeros at all non-zero integer multiples of $2\pi/\Delta$ (Equation 2), but since k and l never assume the value 0 in Equation 3, the autocorrelation function is zero when the lag is not 0 mod L . When the lag is 0 mod L :

$$E \left\{ \exp\left(\frac{j2\pi}{\Delta}(kz_L[n] - lz_L[n+m])\right) \right\} = E \left\{ \exp\left(\frac{j2\pi(k-l)z_L[n]}{\Delta}\right) \right\} = \delta[k-l].$$

This results in:

$$R_{e_L, e_L}[n, n+m] = \frac{\Delta^2}{2\pi^2} \sum_{k=1}^{\infty} \frac{1}{k^2} \cos\left(\frac{2\pi k}{\Delta}(s[?1] - s[n+m])\right). \quad (4)$$

Setting $m = 0$ in (4) and evaluating the resulting summation [9, page 7] yields the power in $e_L[n]$: $R_{e_L, e_L}[n, n] = \Delta^2/12$. From Equation 4, $e_L[n]$ is a cycle-stationary process because $s[n]$ has a finite period, N . Using the results of Ljung [10], spectral information is obtained when Equation 4 is averaged over time. Note that when the lag, m , is not only an integer multiple of L , the period of the dither, but also an integer multiple of N , the autocorrelation function equals $\Delta^2/12$, independently of n . The smallest non-zero lag that satisfies these two conditions is the least common multiple of L and N , denoted by qL where q is an integer. Therefore, the period of the time-averaged autocorrelation function, $\bar{R}_{e_L}[m] = \text{Avg}_n(R_{e_L, e_L}[n, n+m])$, is at least L and at most qL . Let the period equal cL , where c is an integer, $1 \leq c \leq q$. The function $\bar{R}_{e_L}[m]$ can be expressed as a sum of cL weighted complex exponentials:

$$\bar{R}_{e_L}[m] = \sum_{l=0}^{cL-1} p_l \exp\left(\frac{j2\pi ml}{cL}\right), \quad m = \dots, -1, 0, 1, 2, \dots$$

where

$$p_l = \frac{1}{cL} \sum_{m=0}^{cL-1} \bar{R}_{e_L}[m] \exp\left(\frac{j2\pi ml}{cL}\right) = \frac{1}{cL} \sum_{n=\tilde{m}^2}^c \bar{R}_{e_L}[nL] \exp\left(\frac{j2\pi nl}{cc}\right).$$

The last equality is true since the autocorrelation function in Equation 3 and its time-average, $\bar{R}_{e_L}[m]$ are zero for lags not equal to integer multiples of L . The weights, $p_l, l = 0, 1, \dots, cL - 1$, are the power magnitudes of the spur. Since $\bar{R}_{e_L}[m] \leq A^2/12$, the spur power can be bounded: $p_l \leq A^2/(12cL) \leq \Delta^2/(12L)$. Equality is achieved when the period of the time-averaged autocorrelation function is exactly L , the period of the dither.

As $L \rightarrow \infty$, the spacing between spurs goes to zero in the spectra of both $e_L[n]$ and $z_L[n]$. The power in an individual spur goes to zero, but the density (power per unit of frequency) tends to a constant. Thus, ideal white noise behavior is approached. While $z_L[n]$ and $e_L[n]$ are correlated in general, the worst case spur power scenario coherently adds the power spectra from both processes. For this reason, L should be chosen to satisfy $A^2/(12L) < P_{\max}$, where P_{\max} is the maximum acceptable spur power. When constructing a dither signal as the sum of $M \geq 1$ independent, uniform variates the noise autocorrelation becomes $R_{z_L, z_L}[m] = (JA^2/12) \delta[m \bmod L]$. The analysis follows closely to that for $M = 1$, and L should be chosen to satisfy $(A^2 + 1)\Delta^2/(12L) < P_{\max}$.

As in the previous section, since the desired signal has finite word length, it is equivalent to round or truncate the dither signal to an appropriate word length. An implementation using a PN generator is an approximation to such a truncated periodic noise source which produces a periodic sequence of discretely and evenly distributed random numbers.

7 Phase Dithering

In this section, phase dithering is analyzed using a continuous, zero-mean, wide-sense stationary sequence. As described in Section 5 on amplitude dithering, an evenly distributed

discrete random sequence is equivalent to continuous uniform dithering when the initial phase word is quantized to a finite number of bits.

Let the digital sinusoid to be generated be:

$$x[n] = \cos(2\pi(\phi[n] + c[n])) \quad (5)$$

so that the desired phase is $\phi[n]$ as defined in Section 4. The total quantization noise is $c[n] = c_P[n] + z[n]$, the sum of the dither signal and the quantizer error. Using small angle approximations:

$$x[n] = \cos(2\pi fn + \Phi) - 2\pi c[n] \sin(2\pi fn + \Phi) + O((\max(c[n]))^2)$$

The total quantization noise will be examined by considering the first two terms above, and then the second-order, $O((\max(c[n]))^2)$, effect.

7.1 First Order Analysis

Since the quantization error after dithering is independent of the input signal [2] $c[n]$ is uncorrelated with the desired sinusoids. Without loss of generality, and for ease of notation, let us shift the uniformly distributed dither random variate range to $[0, \Delta_P)$. The total phase quantization noise $c[n]$ will be $c[n] = -p[n]\Delta_P$ with probability $(1 - p[n])$, and $c[n] = (1 - p[n])\Delta_P$ with probability $p[n]$. The value $p[n]$ is the distance from the initial high-precision phase value, $\phi[n]$, to the nearest greater quantized value normalized by the phase quantization step size Δ_P . The value of the probability sequence $p[n]$ varies periodically, since $p[n] = \phi[n] \bmod \Delta_P$, and $\phi[n]$ is periodic; however, at all sample times n the first moment of the total phase quantization noise, $E\{c[n]\}$, is zero.

Information about the spurs and noise in the power spectrum of $x[n]$ is obtained from the autocorrelation function. The autocorrelation of $x[n]$ is:

$$E\{x[n]x[n+m]\} = \cos(2\pi fn + \Phi) \cos(2\pi f(n+m) + \Phi)$$

$$+4\pi^2 \sin(2\pi fn + \Phi) \sin(2\pi f(n+m) + \Phi) E\{c[n]c[n+m]\} + O(\Delta_P^4).$$

Spectral information is obtained by averaging over time [10], resulting in:

$$\hat{R}_{xx}[m] \approx \frac{1}{2} [1 + 4\pi^2 \hat{R}_{cc}[m]] \cos(2\pi fm)$$

where $\hat{R}_{cc}[m] = \text{Avg}_n(E\{c[n]c[n+m]\})$, the time-averaged autocorrelation of the total quantization noise.

The power spectrum of $x[n]$, the Fourier transform of the autocorrelation, is the power spectrum of the desired sinusoid of frequency f plus the total quantization noise amplitude modulated on the desired sinusoid. Note that since $\hat{R}_{cc}[m] = O(\Delta_P^2)$, and $\Delta_P \ll 1$, the modulation index is small.

To a first order, the AM signal produced by phase dithering is clear of spurious harmonics down to the level due to periodicities in the dither sequence. The next section will examine spur performance in more detail, but first it is important to consider the noise power spectral density resulting from the phase dithering process.

Recall that for any fixed time n , the probability distribution of $c[n]$, a function of $p[n]$, is determined by the input signal, but the outcome of $c[n]$ is determined entirely by the outcome of the dither signal $z[n]$. When $z[n]$ and $z[n+m]$ are independent random variables for non-zero lag m , $c[n]$ and $c[n+m]$ are also independent for $m \neq 0$, and hence $c[n]$ is spectrally white. In this case, the autocorrelation becomes:

$$\hat{R}_{xx}[m] \approx \frac{1}{2} \cos(2\pi fm) + 2\pi^2 \delta[m] \text{Var}(c)$$

where $\text{Var}(c)$ is the time-averaged variance of the total quantization noise. The resulting signal-to-noise ratio (SNR) is approximately $1/(4\pi^2 \text{Var}(c))$.

When the dither signal is constructed from one uniform $[-\Delta_P/2, \Delta_P/2]$ random variate, the error $c[n]$ is bounded between $-\Delta_P$ to Δ_P where $\Delta_P = 2^b$ and b is the number of bits

in the phase representation after the word-length is reduced. The number of bits, b , must be large enough to satisfy the small angle assumption earlier, e.g., $b \geq 4$. The c-averaged variance of $c[n]$ is less than or equal to $2^{-2b}/4$, and the SNR is $2^{2b}/\pi^2 = 6.02b - 9.94\text{dB}$.

Since the sinusoid generated is a real signal, the signal power in the SNR will be equally divided between positive and negative frequency components. If the sinusoid is the result of a discrete-time random process with sampling frequency f_s , then the resulting noise power spectral density (NPSD) will be given by:

$$\begin{aligned} \text{NPSD} &\approx - [\text{SNR}/2 + 10 \log_{10}(f_s/2)] \text{ dBc/Hz} \\ &\leq 6.93 - 10 \log_{10}(f_s/2) - 6.02b \text{ dBc/Hz}. \end{aligned} \quad (6)$$

Table 1 gives noise power spectral densities as a function of the number of bits per cycle, b , at a 160 MHz sampling rate, calculated according to the above formula.

7.2 Second Order Analysis: Residual Spurs

For a worst-case analysis of second order effects, expand the initial cosine from Equation 5 by the sum of angles formula:

$$x[n] = -\cos(2\pi c[n]) \cos(2\pi f n + \Phi) - \sin(2\pi c[n]) \sin(2\pi f n + \Phi).$$

Information about the spurs in the power spectrum of $x[n]$ is obtained from the autocorrelation function at nonzero lags. When the dither sequence, $z[n]$, is a sequence of i.i. d. variates, the autocorrelation function for $x[n]$, with lag m not equal to zero, is:

$$R_{xx}[n, n+m] = E\{x[n]x[n+m]\} = E\{x[n]\}E\{x[n+m]\}.$$

The expected value of $x[n]$ is a deterministic function of time. From the above expression, it follows that spectral information about the random process $x[n]$, with the exception of noise

floor information, is contained in $E\{x[n]\}$, which we call the “expected waveform”. Since $c[n]$ is zero mean at all sample times the expected waveform reduces to:

$$E\{x[n]\} = (1 - 2\pi^2 E\{c^2[n]\}) \cos(2\pi fn + \Phi) + O(\Delta_p^3).$$

The form of the expected waveform clearly shows that the spurious content of the signal will be derived from the dependence of second and higher-order moments of the quantization noise. It is this fundamental principle that will ultimately lead to the -12 dBc per phase bit behavior for uniformly phase-dithered sinusoid generation.

It remains to consider the second moment of the total phase quantization noise, $E\{c^2[n]\}$, which we evaluate by using the probability sequence, $p[n]$, from the previous section as $E\{c^2[n]\} = \Delta_p^2(p[n] - p^2[n])$. As shown in the appendix, in the worst case this sequence produces a spur whose frequency is the reflection on the desired signal across $1/4$ the sampling rate, and this worst case is achieved for a large class of frequencies. The expected waveform is:

$$\begin{aligned} E\{x[n]\} &= ((1 - \pi^2 \Delta_p^2/4) + (\pi^2 \Delta_p^2/4) \cos(\pi n)) \cos(2\pi fn + \Phi) + O(\Delta_p^3) \\ &= (1 - \pi^2 \Delta_p^2/4) \cos(2\pi fn + \Phi) + (\pi^2 \Delta_p^2/4) \cos((2\pi f + \pi)n + \Phi) + O(\Delta_p^3), \end{aligned}$$

clearly showing the desired signal and spur components. Thus, dropping the $O(\Delta_p^3)$ term, a -18 dB per bit power behavior, the worst-case spur level relative to the desired signal after truncating to b bits is (LSp, \$’it):

$$SpSR \approx \frac{\pi^4 \Delta_p^4/16}{(1 - \pi^2 \Delta_p^2/4)^2} \approx \frac{\pi^4 \Delta_p^4}{16} = 7.84 - 12.04b \text{ dBc}.$$

In summary, if b bits of phase are output to a look-up table, and B bits of phase ($B > b$) are used prior to truncation, then the addition of an appropriate dithering signal using $(B - b)$ bits will allow the word length reduction without introducing spurs governed by the usual $-6b$ dBc behavior. If a single random variate is added as a dither signal (first-order

dithering), the spur suppression is accelerated to 12 dB per bit of phase representation. Since the table size is only effected linearly by the number of bits in a table entry, rather than exponentially as it is by the number of phase bits, the amplitude word length is of secondary importance to the phase word length, especially in all-digital systems. For example, -90 dBc spur performance would nominally require $b = 16$ bits of phase and a 65,536 entry table. With first-order dithering, this level of performance requires only $b \geq 8.1$ bits of phase per cycle in the look-up table addressing. Worst case spur performance of -100.5 dBc is achieved with 9 bits, a 512 entry table at most, and, at a 160 MHz sampling rate, Table 1 shows that with these realistic system parameters, the noise power spectral density is at a low -126 dBc/Hz.

8 Accelerated Spur Suppression

Further analysis [11] based on an extension of results by Gray [12] indicates that the phase spur suppression rate can be increased in steps of 6 dBc per bit by adding multiple uniform random deviates to the phase value prior to truncation. The addition of M uniform random deviates produces a dither signal with $M+1$ -order zeroes in its characteristic function) thus making the M th moment of the quantization error independent of the input sequence [12].

An example of this technique providing 18 dBc per phase bit spur performance is shown in Fig. 6. This technique involves adding two $(B - b)$ -bit uniform deviates to produce a $(B - b + 1)$ bit dither signal, which achieves the accelerated spur reduction due to second-order zeroes in the dither characteristic function. Simulation results for when two uniform variates are added to the phase are presented in the next section. A straightforward extension of this technique to polynomial series allows spur-reduced synthesis of periodic, digital signals with arbitrary waveforms.

9 Simulation Results

Simulations were performed to validate the results of this analysis. These results were obtained using 8192-point unwindowed FFTs, and the synthesized frequencies were chosen to represent worst-case amplitude and phase spur performance. In each of the figures, ten power spectra were averaged to better show the spurious content of the signals. Fig. 7 shows the power spectrum of a sine wave of one-eighth the sampling frequency truncated to 8 bits of amplitude without dithering. Fig. 8 shows the same spectrum with a sixteen-bit sinusoid amplitude dithered with one uniform variate prior to truncation to 8 bits. Note that the spurs have been eliminated to the levels consistent with those imposed by the initial sixteen bit quantization.

Fig. 9 shows the spectrum of a 5-bit phase-truncated sinusoid with high-precision amplitude values. A worst-case example of first-order phase dithering is shown in Fig. 10. The measured noise power spectral density in Fig. 10 is -62.3 dBc per 1° bin, giving a noise density of $-23.2 - 10 \log(f_s/2)$ dBc, in agreement with the upper bound derived in Equation 6. The spur level is -52.3 dBc in the first-order dithered Fig. 10.

Fig. 11 shows the same example using second-order ($M = 2$) dithering using the sum of two uniform deviates. While the spectrum in Fig. 10 shows the residual spurs at -12 dBc per bit due to second-order effects, Fig. 11 shows no visible spurs, indicating better than -63 dBc spurious performance. Additional simulations involving Megapoint FFTs and not represented by figures confirm the -18 dBc per bit performance of the second-order phase-dithered system.

Finally, Fig. 12 shows a worst-case result for first-order phase dithering together with first-order amplitude dithering. The amplitude samples are truncated to 8 bits, as are the phase samples. Note that the spurs are not visible in the spectrum; however, close analysis has demonstrated that they are present at the -88 dBc level expected due to second-order

effects.

10 A System Design Example

The block diagram of a direct digital frequency synthesizer based on the techniques presented here is shown in Fig. 13. The following system would perform at a sampling rate of 160 MHz, producing 8-bit digital sinusoids spur-free to -90 dBc with better than -120 dBc/Hz noise power spectral density. The system parameters are as follows:

Phase bits are in unsigned fractional cycle representation with:

phase accumulator word-length determined by frequency resolution, and

> 16 bits prior to addition of 1 uniform phase dither variate, with ≥ 9 bits after dither addition and truncation;

Amplitude look-up-table with:

$> 2^7 = 128$ entries (using quadrant symmetries) of ≥ 16 bits each normalized so that the sinusoid amplitude equals 512 16-bit quantization steps less than the full-scale value;

Linear feedback shift register PN generator with ≥ 16 lags producing one 8-bit amplitude dither variate, and

One LFSR PN generator with ≥ 18 lags for generation of the 7-bit phase dither variate.

11 Conclusion

A digital dithering approach to spur reduction in the generation of digital sinusoids has been presented. A class of periodic dithering signals has been analyzed because of its similarity

to LFSR PN generators.

The advantage gained in amplitude dithering provides for spur performance at the original longer word length in an ideal system when the digital dithering signal is white noise distributed evenly, not uniformly, over one quantization interval. The reduced word length allows the use of less complicated multipliers and narrower data paths in purely digital applications. If the waveform is ultimately converted to an analog value, the reduced word length allows the use of fast, coarse-resolution, highly-linear digital-to-analog converters (DACs) to obtain sinusoids or other periodic waveforms whose spectral purity is limited by the DAC linearity, not its resolution. These results suggest that coarsely quantized, highly-linear techniques for digital-to-analog conversion such as delta-sigma modulation would be useful in direct digital frequency synthesis of analog waveforms.

The advantage gained in the proposed method of phase dithering provides for an acceleration beyond the normal 6 dB per bit spur reduction to a $6(M+1)$ dB per bit spur reduction when the dithering signal consists of M uniform variates. Often the most convenient way to generate a periodic waveform is by table look-up with a phase index. Since the size of a look-up table is exponentially related to the number of phase bits, this can provide a dramatic reduction in the complexity of NCO's, frequency synthesizers, and other periodic waveform generators.

The advantages of dithering come at the expense of an increased noise content in the resulting waveform. However, the noise energy is spread throughout the sampling bandwidth. In high bandwidth applications, dithering imposes modest system degradation. It has been shown that high performance synthesizers with dramatically reduced complexity can be designed using the dithering method, without, resulting in high noise power spectral density levels.

12 Appendix: Worst-case phase spur analysis

Since the sequence $p[n]$ is bounded between 0 and 1, the function $u[n] = p[n] - p^2[n]$ is bounded between 0 and 1/4, with its maximum value of 1/4 at $p[n] = 1/2$. Therefore, $u[n]$ must have some non-zero 11~ (average) component. Any remaining components can be periodic in the worst case. Since all non-linear operations have been performed, conservation of power (energy) arguments can be used to determine the total non-DC error power. The total power in the DC component of $u[n]$ is equal to the square of the average value of $u[n]$. Similarly, the total power in $u[n]$ is equal to the average value of $u^2[n]$. Thus, the power remaining for time-varying components of $u[n]$ is :

$$\text{Avg}(u^2[n]) - (\text{Avg}(u[n]))^2 = \text{Avg}((u[n] - \text{Avg}(u[n]))^2).$$

This value is maximized by maximizing the dispersion of the samples about the mean. When the sample values are bounded, this maximization is achieved by placing half of the samples at each bound, so that the mean is equidistant from each bound. Since $0 \leq u[n] \leq 1/4$, the maximum power present in harmonic components is 1/64.

‘ 1 ’ 0 find a sequence achieving this bound, the sequence $E\{c^2[n]\}$ must be examined in more detail. The previous paragraph indicates that the bound will be achieved when half the samples of $E\{c^2[n]\}$ are 0 and half are $\Delta^2/4$. Since $c^2[n]$ is non-negative, $E\{c^2[n]\} = 0$ implies that $c[n] = 0$. Note that the difference between $c[n]$ and $c[n+1]$ is the phase increment modulo the quantization step size. If, for any n and $n+1$, $c[n] = c[n+1] = 0$, the phase increment can be exactly expressed in the new quantization step. By induction, $c[n]$ will be zero for all n if any two adjacent values $E\{c^2[n]\}$ and $E\{c^2[n+1]\}$ are both zero. The only possible sequence $E\{c^2[n]\}$ achieving the worst case is therefore 0, 1/4, 0, 1/4, 0, 1/4..., and thus, $u[n] = 1/8 - (1/8) \cos(\pi n)$. This sequence has a single sinusoidal component at the Nyquist frequency, half the sampling rate.

13 Acknowledgements

The research described in this paper was carried out at the Department of Electrical Engineering and at the Jet Propulsion Laboratory of the California Institute of Technology, under a contract with the National Aeronautics and Space Administration. Michael J. Planagan was supported in part by a National Science Foundation Fellowship.

References

1. A.B. Sripad and D.L. Snyder, "A necessary and sufficient condition for quantization errors to be uniform and white," *IEEE Transactions on Acoustics Speech and Signal Processing*, vol. ASSP-25, pp. 442-448, October 1977.
2. L. Schuchman, "Dither signals and their effects on quantization noise," *IEEE Transactions on Communication Technology*, vol. COM-12, pp. 162-165, December 1964.
3. N.S. Jayant and L.R. Rabiner, "The application of dither to the quantization of speech signals," *Bell System Technical Journal*, vol. 51, pp. 1293-1304, 1972.
4. S.C. Jasper, "Frequency Resolution in a Digital oscillator," U.S. Patent Number 4,652,832, March 24, 1987.
5. P. O'Leary and F. Maloberti, "A Direct-Digital Synthesizer with Improved Spectral Performance," *IEEE Transactions on Communications*, vol. COM-39, pp. 1046-1048, July 1991.
6. J. Tierney, C. Rader, B. Gold, "A Digital Frequency Synthesizer," *IEEE Transactions on Audio and Electroacoustics*, vol. AU-19, No. 1, March 1971, pp. 48-57.

7. T. Nicholas and H. Samueli, "An Analysis of the Output Spectrum of Direct Digital Frequency Synthesizers in the Presence of Phase Accumulator Truncation," *41st Annual Frequency Control Symposium*, USFIRACOM, Ft. Monmouth, NJ, pp. 495-502, 1987.
8. L.E. Brennan and J.S. Reed, "Quantization Noise in Digital Moving Target Indication Systems," *IEEE Transactions on Aerospace and Electronic Systems*, vol. AES-2, pp. 655-658, November 1966.
9. J.S. Gradshteyn and I.M. Ryzhik, 'Table of Integrals, Series, and Products, New York: Academic Press, corrected and enlarged edition, 1980.
10. T. Ljung, System Identification: Theory for the User, Prentice-Hall, Englewood Cliffs, NJ, 1987.
11. M.J. Flanagan and G.A. Zimmerman, "Spur-Reduced Digital Sinusoid Generation Using Higher-order Phase Dithering," Conference Record: Papers Presented at the 27th Annual Asilomar Conference on Signals, Systems and Computers, Pacific Grove, CA, IEEE Computer Society Press, November 1993, in press.
12. R.M. Gray and T.G. Stockham, Jr., "Dithered Quantizers," *IEEE Transactions on Information Theory*, vol. IT-39, pp. 805-812, May 1993.

Figures

1. Spur generation in conventional digital sinusoid generation
2. Input/output relation of a midtread quantizer
3. Conceptual waveform generator model

4. Two-step waveform generator model
5. Uniform dithered quantizer
6. System for 18 dBc per phase bit spur reduction
7. Power spectrum of 8 sample/cycle sine wave without dithering. (8 bit amplitude quantization)
8. Power spectrum of 8 sample/cycle sine wave with amplitude dithering. (8 bit amplitude quantization)
9. Power spectrum of 5-bit phase-truncated sine wave without phase dithering. (high-precision amplitude)
10. Power spectrum of 5-bit phase-truncated sine wave with first-order phase dithering. (high-precision amplitude)
11. Power spectrum of 5-bit phase-truncated sine wave with second-order phase dithering. (high-precision amplitude)
12. Worst-case power spectrum of sinusoid with first-order phase dithering and amplitude dithering (8 bits each).
13. Block diagram of spur-reduced direct digital frequency synthesizer.

Table I: Noise Power Spectral Densities **for 160 MHz Sampling Rate**

b (bits/cycle) Noise Power Spectral Density

5	-102.20 dBc/Hz
6	-108.22 dBc/Hz
7	-114.24 dBc/Hz
8	-120.26 dBc/Hz
9	-126.28 dBc/Hz
10	-132.30 dBc/Hz
11	-138.32 dBc/Hz
12	-144.35 dBc/Hz

Fig. 1: Spur generation in conventional digital sinusoid generation

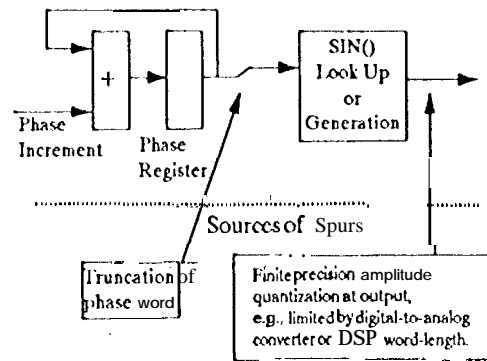


Fig. 2: Input/output relation of a midtread quantizer

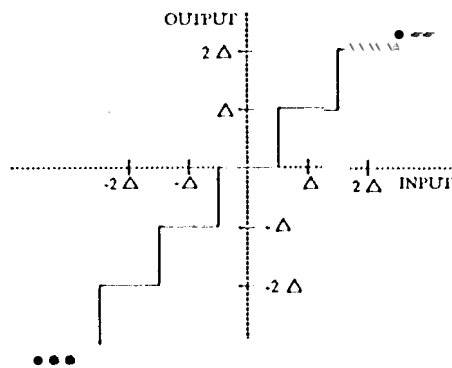


Fig. 3: Conceptual waveform generator model

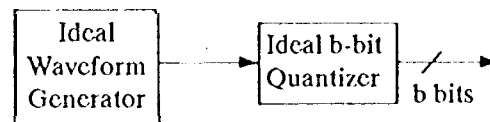


Fig. 4: Two-step waveform generator model

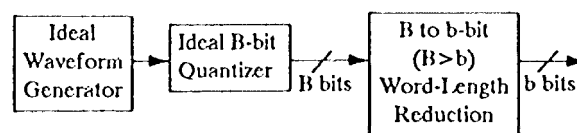


Fig. 5: Uniform dithered quantizer

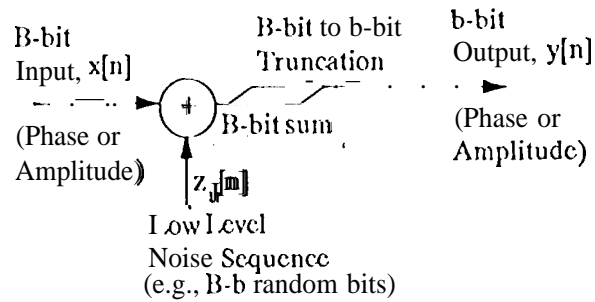


Fig. 6: System for 18 dBc per phase bit spur reduction

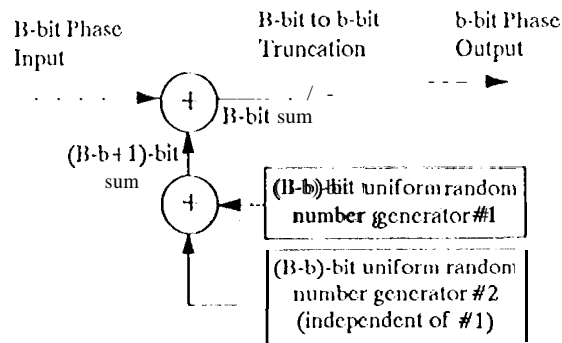


Fig. 7: Power spectrum of 8 sample/cycle sine wave with dithering (8 bit amplitude quantization)

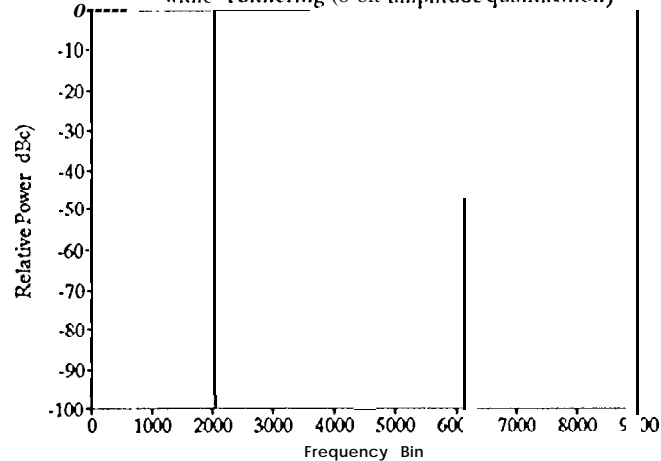


Fig.8: Power spectrum of 8 sample./cycle sine wave with amplitude dithering (8 bit amplitude quantization)

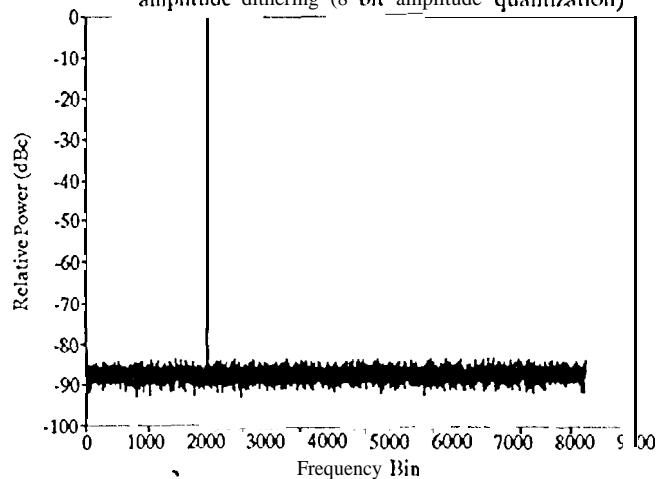


Fig. 9: Power spectrum of 5-bit phase-truncated sine wave without phase dithering (high-precision amplitude)

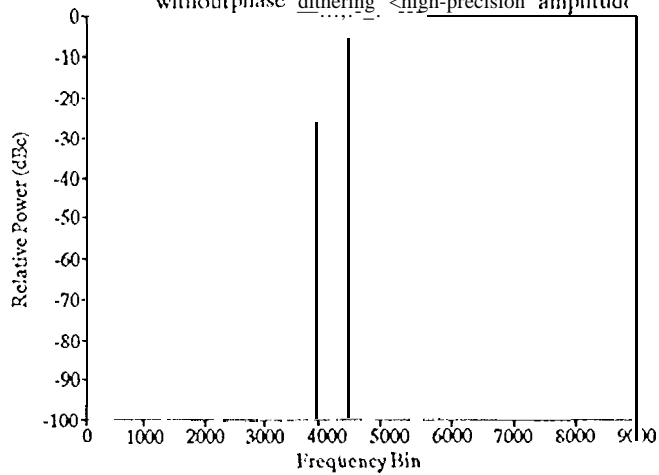


Fig.10: Power spectrum of 5-bit phase-truncated sine wave with first-order phase dithering (high-precision amplitude)

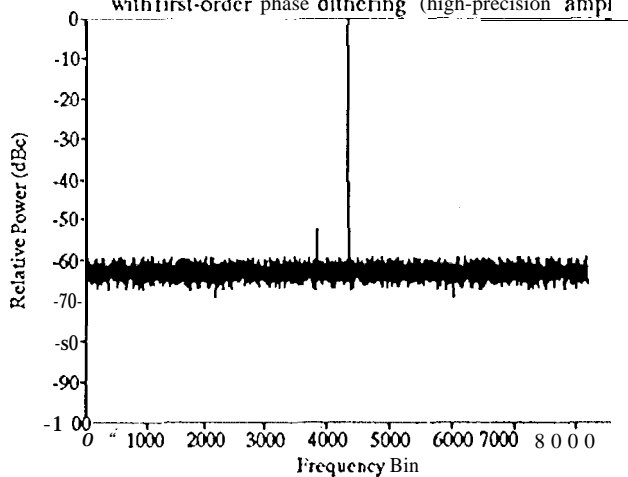


Fig. 11: Powerspectrum of 5-bit phase-truncated sine wave with second-order phase dithering (high-precision amplitude)

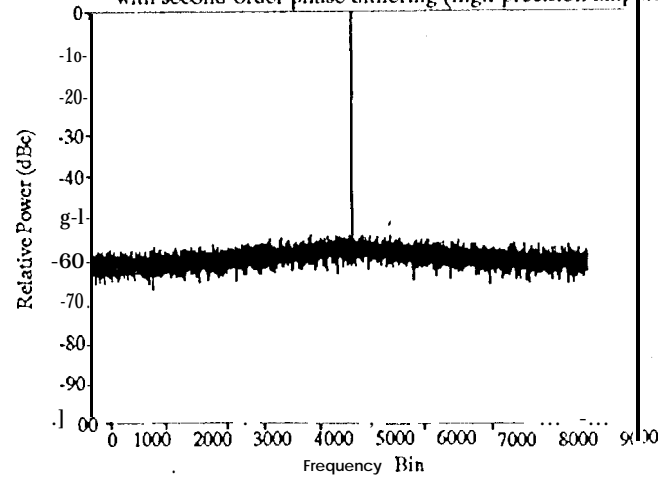


Fig. 12: Worst-case powerspectrum of sinusoid with first-order phase dithering and amplitude dithering (bits each).

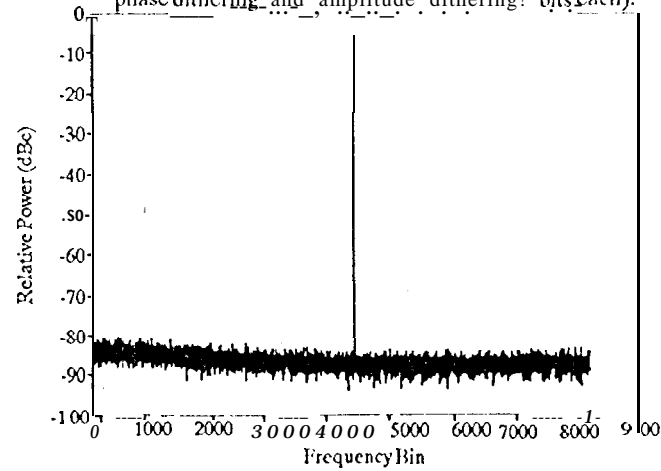


Fig. 13: Block diagram of spur-reduced direct digital frequency synthesizer

